arXiv:1511.08190v1 [stat.ME] 25 Nov 2015

A latent trawl process model for extreme values

Ragnhild C. Noven Almut E. D. Veraart Axel Gandy

Department of Mathematics, Imperial College London

Abstract

This paper presents a new model for characterising temporal dependence in exceedances above a threshold. The model is based on the class of trawl processes, which are stationary, infinitely divisible stochastic processes. We review properties of trawl processes in the context of statistical modelling, and introduce a new representation that enables exact simulation for discrete observations. The model for extreme values is constructed by embedding a trawl process in a hierarchical framework, which ensures that the marginal distribution is generalised Pareto, as expected from classical extreme value theory. We also consider a modified version of this model that works with a wider class of generalised Pareto distributions, and has the advantage of separating marginal and temporal dependence properties. The model is illustrated by applications to environmental time series, and it is shown that the model offers considerable flexibility in capturing the dependence structure of extreme value data.

1 Introduction

Modelling dependencies in extreme value data is a topic of growing importance, having applications in a number of fields such as hydrology (de Haan and de Ronde, 1998), oceanography (Coles and Tawn, 1991), financial risk management (Ledford and Tawn, 2003; Embrechts et al., 1997) and environmental science (Davison et al., 2012; Heffernan and Tawn, 2004).

For extreme value data it is particularly valuable to have a model that captures the dependence structure, for two reasons: first, it allows us to borrow strength across multiple observations, which is especially relevant as extreme data are scarce by definition. Second, the impact of extreme events may depend on the occurrence of other events in the past, present or future; for example, the risk of flooding depends heavily on whether rainfall or storm surge events tend to occur close together in time.

Historically, the focus of extreme value analysis has been on modelling the tails of univariate i.i.d. observations. A series of papers, summarised in Leadbetter et al. (1983), expanded the existing theory to hold for stationary sequences under some mild assumptions. There has also been much work on models for multivariate extremes, which have been applied to both time series and observations from multiple sources (Coles and Tawn, 1991; Bortot and Tawn, 1998).

A recurring problem when modelling dependent extremes is the trade-off between the tractability of a given model and its ability to capture the full dependence structure. The class of max-stable processes appears promising due to its flexibility and grounding in extremal theory, and has been used for spatio-temporal modelling (see Huser and Davison, 2014, and references therein). However, the computational burden associated with estimation appears to be a general problem for this class of models. Furthermore, the models fail to represent the typical processes found in environmental applications as they do not allow for realistic temporal dynamics – the time and space aspects are treated separately and given similar dependence structures (this is discussed in Embrechts et al., 2015, who propose an alternative class of max-stable models).

In order to obtain tractable models, strong assumptions are often made about the dependence structure of extremes. It is often the case that a given time series suggests a natural block division (e.g. with blocks corresponding to months or years), in which case the modelling is done on block maxima, which are assumed to be independent. Although this assumption is justified asymptotically, this procedure makes inefficient use of the data, which in the extreme value case are scarce to begin with; furthermore, it does not describe the dependence structure within blocks, which may be of independent interest.

A popular alternative to working with block maxima is to consider exceedances above a threshold, based on the "peaks over threshold" (POT) method used in hydrology (Davison and Smith, 1990), which fits the exceedances to a generalised Pareto distribution. However, this approach requires independent observations, and thus in practice it is common to use a so-called declustering method on the raw data, which tries to identify clusters in the data and use these to extract approximately independent observations of the quantity of interest. Declustering methods are based on a fundamental result in Hsing et al. (1988), which shows that the normalised point process of exceedances converges to a compound Poisson process as the threshold increases, and that the cluster maxima are independent in the limit. Although they are theoretically justified, declustering methods will inevitably introduce additional uncertainty associated with the choice of declustering method and associated parameters (though see Ferro and Segers, 2003, for an automated selection procedure), and as with the block maxima approach they do not use all of the available data.

Among models that try to model the dependence explicitly, a common assumption is that a given times series of extremes has the Markov property. This assumption motivates the approach used by Smith et al. (1997), who consider a Markov chain with transition probabilities derived from bivariate extremal dependence models. Ledford and Tawn (1997) consider a more general dependence structure based on an asymptotic expansion of the joint survivor function, which leads to a wider class of bivariate dependence models.

It is clear from the above that the main current approaches to modelling time series of extremes do not allow for a complete description of the temporal dependence structure for processes that are not essentially independent or Markov. One could in theory work with higher order Markov chains and use multivariate extremes models of higher dimensions, but this quickly becomes intractable in practice due to the computational burden and the lack of a finite parameterisation for the joint distributions of such models. A more fundamental issue is the application of multivariate extremes models to time series without any consideration for the time-specific aspects such as causality, which could potentially be used to reduce the dimensionality of the problem.

Motivated by the above concerns, this paper presents a new model for characterising temporal dependence in exceedances above a threshold. The model makes use of the class of trawl processes (Barndorff-Nielsen, 2011), which are stationary, non-Markovian stochastic processes. The class of trawl processes belongs to the more general class of ambit processes, which were originally developed as general spatio-temporal models for studying turbulence (Barndorff-Nielsen and Schmiegel, 2007; Barndorff-Nielsen et al., 2015).

The new model uses the same hierarchical structure as Bortot and Gaetan (2014), which ensures that the marginal distribution of the exceedances is generalised Pareto, which is the proper limiting distribution for exceedances from classical extreme value theory. Our trawl-based model allows for a flexible dependence structure, only imposing some assumptions that have clear physical interpretations (see section 2.4), in contrast with the Markov chain setup used in Bortot and Gaetan (2014).

The contribution of the present paper is threefold: First, we develop a representation for trawl

processes in terms of independent random variables, which enables exact simulation for discrete observations. Due to this representation, the trawl process subclass may be more attractive for certain applications than the more general Ambit or Lévy semistationary processes, which require computationally intensive approximations (Bennedsen et al., 2014; Hedevang and Schmiegel, 2014). Second, we develop a new model for exceedances, which uses trawl processes to obtain a more flexible dependence structure than the model in Bortot and Gaetan (2014), which uses a similar hierarchical setup. Finally, we consider a modified version of the model that enables it to be used for generalised Pareto distributions with a negative shape parameter, and also has the effect of separating the marginal and dependence structures of the model.

This paper is structured as follows. Section 2 starts by reviewing the necessary theoretical background for trawl processes, proceeds to introduce a new representation, and finishes with a discussion of the temporal dependence structure that gives guidance on choosing a trawl process for a particular statistical model.

Section 3 introduces the new model for dependent extremes and a modification for generalised Pareto distributions with a negative shape parameter. Section 4 discusses parameter estimation, develops a measure for the extremal dependence structure that is adapted to the model, and shows examples of applying the model to two different environmental time series.

2 Trawl process background

In this section we introduce trawl processes, which are a class of stationary stochastic processes defined in terms of an underlying Lévy basis (Barndorff-Nielsen, 2011; Barndorff-Nielsen et al., 2015). A trawl process is made up of two components: the Lévy basis, which is a particular random measure defined on a Borel space S, and a fixed subset of S called the trawl set.

We start by introducing the theoretical background for the Lévy basis, and then we define the trawl process, following the expositions in Barndorff-Nielsen (2011) and Barndorff-Nielsen et al. (2015). The remainder of the section then introduces two novel representations of the trawl process: in subsection 2.3 we derive a representation in terms of a collection of independent random variables, which is particularly important for modelling purposes because it enables exact simulation. Subsection 2.4 focuses on the temporal structure of the trawl process, in particular the autocorrelation function, demonstrating that the trawl process can be interpreted as accumulations of random pulses from a causal system.

2.1 Lévy basis definition

Consider a probability space $(\Omega, \mathcal{F}, P)$ and let S be a Borel set in $\mathbb{R}^n$, with the associated Borel σ-algebra $\mathcal{S} = \mathcal{B}(S)$ and Lebesgue measure λ. Let $\mathcal{B}_b(S)$ be the subsets of S with finite Lebesgue measure λ, i.e.

$$\mathcal{B}_b(S) = \{A \in \mathcal{S} : \lambda(A) < \infty\}.$$

A *Lévy basis* is a particular type of random measure, which we define as follows:

Definition 2.1: A *random measure on* $(S, \mathcal{B}(S))$ is a collection of $\mathbb{R}$-valued random variables

$$\{M(A) : A \in \mathcal{B}_b(S)\}$$

such that for any sequence $A_1, A_2, \ldots$ of disjoint elements of $\mathcal{B}_b(S)$ with bounded union, i.e. $\cup_{j=1}^{\infty} A_j \in \mathcal{B}_b(S)$, we have

$$M\left(\bigcup_{j=1}^{\infty} A_j\right) = \sum_{j=1}^{\infty} M(A_j) \quad a.s.$$

Remark 2.2: The above definition only considers sets of bounded Lebesgue measure, this is sufficient, however, as $\mathcal{B}(S)$ is generated by the sets in $\mathcal{B}_b(S)$. As will be seen, the area $\lambda(A)$ of a set A enters into the distribution of the random variable $M(A)$, and thus it is convenient to work with sets of bounded area.

Remark 2.3: This definition of a random measure differs from that in Kallenberg (1976), where random measures are defined as maps from the probability space Ω to the set $\mathcal{M}(S)$ of Radon measures on S. In contrast, the random measure defined above is analogous to a stochastic process, indexed by sets in $\mathcal{B}_b(S)$ as opposed to time points.

To illustrate this difference, consider the set function $\phi_\omega : S \to \mathbb{R}$ defined on $\mathcal{B}_b(S)$ via $\phi_\omega(A) := M(A)(\omega)$. Because this function is allowed to take negative values, it may fail to be an ordinary measure; more importantly, it may also fail to be a signed measure. The reason for this is that the additivity condition in the random measure definition only requires almost sure equality, where the associated null-set may depend on the particular sequence $\{A_i\}$. For example, if we let ϕ_ω be the path of a standard Brownian motion, it will have infinite variation as a function on $\mathbb{R}$, and so in particular will not be a signed measure.

To emphasise that the above definition differs from that of Kallenberg (1976), some authors use the term *random noise* for the random measure M defined above. However, the distinction between random measures and random noise will not be important here, as we will only be concerned with finite-dimensional distributions in the following.

The next definition gives an analogue of the independent increment property of Lévy processes applied to random measures.

Definition 2.4: We say that a random measure M on $(S, \mathcal{S})$ is *independently scattered* if for any sequence $A_1, A_2, \ldots$ of disjoint elements of $\mathcal{B}_b(S)$, the random variables $M(A_1), M(A_2), \ldots$ are independent. The term "completely random measure" is also sometimes used for this property.

Finally, we define infinite divisibility for random measures.

Definition 2.5: A random measure M on $(S, \mathcal{S})$ is called *infinitely divisible* if for each $n \in \mathbb{N}$ there exist n independent, identically distributed random measures $Z_1^n, \ldots, Z_n^n$ such that

$$M \stackrel{d}{=} Z_1^n + \ldots + Z_n^n.$$

In particular, infinite divisibility implies that for any finite collection $A_1, \ldots, A_n$ of elements of $\mathcal{B}_b(S)$, the random vector $(M(A_1), \ldots, M(A_n))$ is infinitely divisible in $\mathbb{R}^n$.

With the above definitions to hand, we are now ready to define a Lévy basis.

Definition 2.6: A *Lévy basis* L on $(S, \mathcal{S})$ is a random measure that is independently scattered and infinitely divisible.

In the following we will work with a so-called *homogeneous* Lévy basis L, meaning that the random measure is stationary, in addition to the standard Lévy basis properties:

Definition 2.7: A random measure on $(S, \mathcal{S})$ is called *stationary* if for any point $s \in S$ and finite collection $A_1, A_2, \ldots, A_n \in \mathcal{B}_b(S)$ such that $A_i + \boldsymbol{s} \subset \mathcal{B}_b(S)$, we have that

$$(M(A_1 + \boldsymbol{s}), M(A_2 + \boldsymbol{s}), \ldots, M(A_n + \boldsymbol{s})) \stackrel{d}{=} (M(A_1), M(A_2), \ldots, M(A_n)).$$

2.1.1 Lévy basis specification

Having defined a Lévy basis, a natural question is how we can construct such an object. This subsection introduces a characterisation that enables us to specify homogeneous Lévy bases explicitly via a generating random variable.

We start with the observation that the random variable $L(A)$ is infinitely divisible for any set $A \in \mathcal{B}_b(S)$, and hence, by the Lévy-Khintchine theorem, it has a characteristic function of form

$$E[\exp\{iuL(A)\}] = \exp\left\{iu\widehat{a}(A) - \frac{1}{2}\widehat{b}(A)u^2 + \int_{\mathbb{R}} (e^{iux} - 1 - iux\mathbb{I}_{[-1,1]}(x))\, n(dx, A)\right\}, \quad (1)$$

for $u \in \mathbb{R}$, where $\widehat{a}(A) \in \mathbb{R}$, $\widehat{b}(A) > 0$ and $n(\cdot, A)$ is a Lévy measure for fixed A, i.e.

$$n(\{0\}, A) = 0, \quad \int (|x|^2 \wedge 1)\, n(dx, A) < \infty.$$

We say that $(\widehat{a}, \widehat{b}, n(\cdot, A))$ is the *characteristic triplet* of $L(A)$.

The precise relationship between L and the characteristic triplet is described by Proposition 2.1 in Rajput and Rosinski (1989), applied to $\mathcal{B}_b(S)$, which is restated below in our notation:

Proposition 2.8:

1. Let L be a Lévy basis with characteristic function given by (1). Then $\widehat{a}$ and $\widehat{b}$ are signed and ordinary measures on $\mathcal{B}_b(S)$, respectively, $n(\cdot, A)$ is a Lévy measure on $\mathbb{R}$ for fixed A, and $n(B, \cdot)$ is a measure on $\mathcal{B}_b(S)$ for fixed $B \in \mathcal{B}(\mathbb{R})$ such that $0 \notin \bar{B}$.
2. Given measures $\widehat{a}, \widehat{b}$ and n as described in part (1), there exists a unique Lévy basis such that (1) holds.

Part 2 of the above proposition implies that we can define Lévy basis by specifying measures $\widehat{a}, \widehat{b}$ and n satisfying the conditions in part 1. This Lévy basis is defined on a probability space $(\Omega, \mathcal{F}, P)$, which accommodates the collection of random variables $L(A)$ for all $A \in \mathcal{B}_b(S)$, with the law of $L(A)$ determined by (1).

When using the definition above, it is convenient to define the measures $\widehat{a}, \widehat{b}$ and n by specifying their densities with respect to the Lebesgue measure on the space $(S, \mathcal{S})$. These densities are denoted a, b and $\nu(\cdot, z)$, where

$$\widehat{a} = \int a(z)\, dz, \quad \widehat{b} = \int b(z)\, dz, \quad n(\cdot, A) = \int_A \nu(\cdot, z)\, dz. \quad (2)$$

It is clear that in order to obtain a homogeneous Lévy measure L, the measures $\widehat{a}, \widehat{b}$ and n must be stationary on $(S, \mathcal{S})$, implying that a, b are constant and $\nu(\cdot, z) = \nu(\cdot)$ is independent of z.

Using the above definitions gives a concrete procedure for specifying a homogeneous Lévy basis. Specifically, consider an infinitely divisible random variable L' with characteristic triplet (γ, σ^2, Π), as defined in the Lévy-Khintchine representation for random variables, and define the densities $a = \gamma$, $b = \sigma^2$ and $\nu = \Pi$. It follows that the resulting measures $\widehat{a}, \widehat{b}$ and n satisfy the Lévy basis conditions in Proposition 2.8. In particular, ν is a Lévy measure, so the relation

$$n(\cdot, A) = \int_A \nu(\cdot)\, dz = \nu(\cdot)\lambda(A),$$

implies that $n(\cdot, A)$ is a Lévy measure for fixed A in $\mathcal{B}_b(S)$. The random variable L' is called the *Lévy seed* of L.

Using the Lévy-Khintchine representation (1), relation (2) and the characteristic triplet $(a, b, \nu(\cdot))$ of L', we obtain the fundamental relation

$$E\left[\exp\{iuL(A)\}\right] = \exp\left\{\lambda(A) \log E\left[\exp\left\{iuL'\right\}\right]\right\}, \quad (3)$$

which shows that the law of $L(A)$ is fully determined by the Lévy seed L' and the Lebesgue measure of the set A.

We now give some examples of homogeneous Lévy bases, where throughout A is an arbitrary set in $\mathcal{B}_b(S)$.

Gaussian basis: Let $\nu(dx) \equiv 0$, so that the Lévy seed L' has characteristic triplet $(a, b, 0)$. Then L is a Gaussian basis, with

$$L' \sim N(a, b), \quad L(A) \sim N\left(a\lambda(A), b\lambda(A)\right).$$

In particular, if we let S be the real line, set $a = 0$, $b = \sigma^2$ and define a process by $Y_t = L((0, t])$, then Y has independent increments and $Y_t \sim N(0, \sigma^2 t)$, so we recover the standard Brownian motion.

Poisson basis: Let $a = b = 0$ and $\nu(dx) = \gamma\delta_1(x)$, where δ_1 is the Dirac delta measure having a point mass at 1. Then

$$L' \sim \text{Poisson}(\gamma), \quad L(A) \sim \text{Poisson}(\gamma\lambda(A)),$$

so L is a Poisson Lévy basis. Working on the real line and defining $Y_t = L((0, t])$ gives the familiar Poisson process with intensity γ.

Gamma basis: Let $a = b = 0$ and $\nu(dx) = \alpha x^{-1} e^{-\beta x}$. Then

$$L'(A) \sim \text{Gamma}(\alpha, \beta), \quad L(A) \sim \text{Gamma}\left(\alpha\lambda(A), \beta\right),$$

so L is a Gamma Lévy basis.

These examples illustrate how the area $\lambda(A)$ scales the cumulants linearly, as follows from the characteristic function relation (3).

2.2 Trawl process definition

In the following we work with the space $S = \mathbb{R}^d \times \mathbb{R}$, with associated σ-algebra $\mathcal{B}(S)$, and the collection $\mathcal{B}_b(S)$ of sets with bounded Lebesgue measure $\lambda(\cdot)$.

Definition 2.9: A trawl process Y_t is generated by taking a homogeneous Lévy basis L and evaluating it over a set A called the trawl. More specifically, let A be any set in $\mathcal{B}_b(\mathbb{R}^d \times \mathbb{R})$, and define a collection of trawl sets $\{A_t\}$ by shifting A along the $\mathbb{R}$-axis corresponding to the last coordinate:

$$A_t = A + (\mathbf{0}, t) := \{(a_1, \ldots, a_d, a_{d+1} + t) : (a_1, \ldots, a_{d+1}) \in A\}.$$

The trawl process Y_t is then defined as

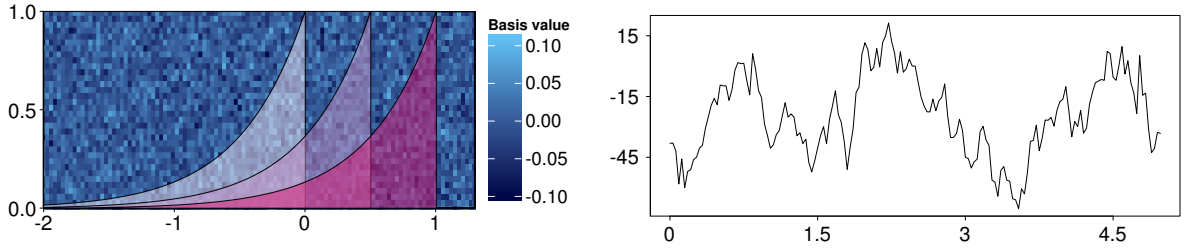
$$Y_t = L(A_t). \tag{4}$$

Remark 2.10: The trawl process definition can be written as a stochastic integral, which will become useful for calculations in the following. Specifically, we write

$$Y_t = \int_{\mathbb{R}^d \times \mathbb{R}} \mathbb{I}_{A_t}(\xi, s) L(d\xi, ds) = \int_{\mathbb{R}^d \times \mathbb{R}} \mathbb{I}_A(\xi, s - t) L(d\xi, ds),$$

where points in $\mathbb{R}^{d+1}$ are denoted (ξ, s) for $\xi \in \mathbb{R}^d, s \in \mathbb{R}$, so the last component corresponds to the time axis. The general stochastic integral is defined in Rajput and Rosinski (1989) and reviewed in Barndorff-Nielsen et al. (2015); it uses a construction based on simple functions, which is a natural generalisation of the trawl process definition.

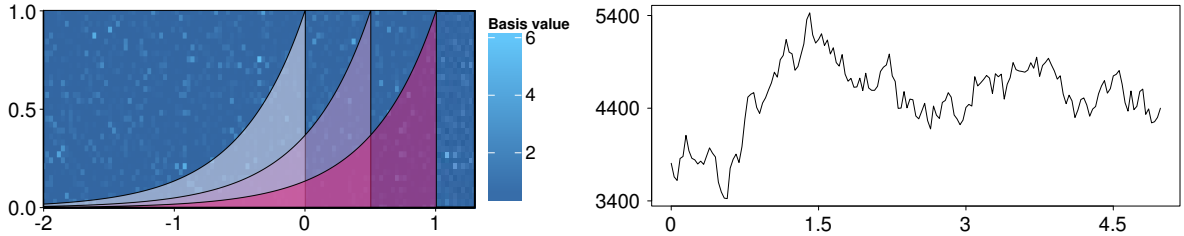
The stochastic integral definition implies that a trawl process is also a so-called mixed moving-average process; it was shown in Fuchs and Stelzer (2013) that such processes are mixing, so it follows that trawl processes are mixing and ergodic.



(a) Discretised Gaussian Lévy basis and trawl sets.

(b) Gaussian trawl process path.

Figure 1: Gaussian Lévy basis with associated trawl process.



(a) Discretised Gamma Lévy basis and trawl sets.

(b) Gamma trawl process path.

Figure 2: Gamma Lévy basis with associated trawl process.

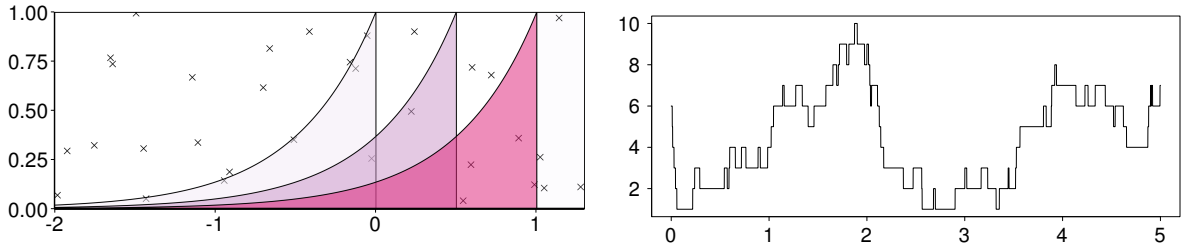
Figures 1 and 2 illustrate the trawl process definition for $d = 1$, showing Gaussian and Gamma Lévy bases, respectively, with a trawl set superimposed at different time points, and the resulting paths of the process. The images show a discretisation of the actual Lévy basis, which is obtained by taking a fine grid covering a sufficiently large subset of $\mathbb{R} \times [0, 1]$, evaluating the random measure L over each of the grid squares and simulating the resulting random variables. The trawl set used in these figures is defined as the area between the the s-axis and the function

$$d(s) = \exp(-2(t - s)), \quad s \in (-\infty, t], \tag{5}$$

this is called an *exponential trawl*. Figure 3 shows a Poisson Lévy basis, represented exactly as points in $\mathbb{R} \times [0, 1]$, with the same trawl sets superimposed.

2.2.1 Distributional properties of trawl processes

The marginal distribution of the trawl process value $Y_t = L(A_t)$ is specified via the characteristic function in equation (3), where L is generated by the Lévy seed L' as described in section 2.1.1. From this equation it follows that Y_t only depends on the trawl set A_t through its size, i.e. the



(a) Poisson Lévy basis and trawl sets.

(b) Poisson trawl process path.

Figure 3: Poisson Lévy basis with associated trawl process.

Lebesgue measure $\lambda(A_t)$. By construction, the shape of the trawl set is constant over time, hence $\lambda(A_t) = \lambda(A)$ for all t, resulting in a stationary process.

The following proposition shows that when using an appropriate normalisation for the cumulant of the Lévy seed L', the particular choice of trawl set does not affect the trawl process marginally.

Proposition 2.11: For any infinitely divisible distribution $\mathcal{D}$ and set $A \in B_b(\mathbb{R}^d \times \mathbb{R})$, there exists a trawl process having trawl set A and marginal law $\mathcal{D}$.

Proof. Consider a random variable $X \sim \mathcal{D}$, and define a Lévy seed L' via its cumulant function[1]

$$C\left(u, L'\right) := \log E\left(\exp\{iuL'\}\right) = \frac{C\left(u, X\right)}{\lambda(A)}.$$

Using (3) then gives

$$E\left[\exp\{iuY_t\}\right] = E\left[\exp\{iuL(A_t)\}\right] = \exp\{\lambda(A_t)C\left(u, L'\right)\} = \exp\{C\left(u, X\right))\},$$

such that $Y_t \sim \mathcal{D}$ for any $t \geq 0$. □

As will be shown in Section 2.4, the autocorrelation structure of the trawl process is fully characterised by the shape of the trawl set. Thus the temporal and marginal properties of the trawl process can be separated, which is a useful property for statistical applications.

We now turn to the joint distribution of trawl process observations, which is specified by the following lemma:

Lemma 2.12: Let $Y_t = L(A_t)$ be a trawl process, then the joint characteristic function of $Y_{t_1}, \ldots, Y_{t_k}$ is given by

$$E\left[\exp\left\{i\sum_{m=1}^{k} u_m Y_{t_m}\right\}\right] = \exp\left\{\int_{\mathbb{R}^d\times\mathbb{R}} C\left(\sum_{m=1}^{k} u_m \mathbb{I}_{A_{t_m}}(y),\, L'\right) dy\right\},$$

Proof. Proposition 5 in Barndorff-Nielsen et al. (2015) gives the cumulant function of the quantity $\mu(Y) = \int_{\mathbb{R}} Y_t\, \mu(dt)$ for any measure μ, assuming this integral exists a.s. and is finite. In particular, they show that

$$C(\theta, \mu(Y)) = \int_{\mathbb{R}^d\times\mathbb{R}} C\left(\theta\left(\int_{\mathbb{R}} \mathbb{I}_A(\xi, s-t)\, \mu(dt)\right), L'\right) d\xi ds.$$

The result follows from this expression when using the measure $\mu(dt) = \sum_{m=1}^{k} u_m \delta_{t_m}(dt)$, with δ being the Dirac delta measure. □

The transition from the marginal to the joint distribution adds complexity to the expression for the characteristic function, as we can no longer reduce the integral to the area of a trawl set. In particular, the presence of several discontinuities in the argument of the cumulant function means that numerical approximations of the integral will usually give poor results.

It is possible to simplify the expression for the characteristic function, by noting that the argument of the integrand $C(\cdot, L')$ is the sum of contributions u_m from all trawl sets containing the point y; hence $C(\cdot, L')$ can only take finitely many different values as y ranges over $\mathbb{R}^d \times \mathbb{R}$, assuming the total number of observations k is fixed. This observation forms the basis of the following subsection, where we use the set of possible values of the cumulant function to define a corresponding set of random variables. This leads to a useful representation of the trawl process that also gives a simplified expression for the joint characteristic function.

[1] The cumulant function of a random variable is the distinguished logarithm of its characteristic function, as defined in Sato (1999, p. 33).

2.3 Partition representation

In this subsection we develop a novel representation of the trawl process in terms of independent random variables, which provides a method for simulating exactly from a general trawl process. We will also consider the particular case of a monotonic trawl, where the representation can be simplified further.

Consider the case where we have k observations $Y_{t_1}, \ldots, Y_{t_k}$ of a trawl process, as defined in (4), with k corresponding trawl sets $A_{t_1}, \ldots, A_{t_k}$. As described in the following definition, these trawl sets are generated by a particular collection of subsets.

Definition 2.13: A *slice partition* of the trawl sets $A_{t_1}, \ldots, A_{t_k}$ is a collection of subsets such that the following conditions hold:

1. Each trawl set A_j is a union of subsets in the partition.
2. The intersection of any number of trawl sets and trawl set complements is a union of subsets in the partition.

In the general case we construct a slice partition by taking all intersections of l trawl sets for $l = 1, \ldots, k$, and using inclusion-exclusion. For example, the slice partition for $k = 2$ consists of three subsets: $\{A_1 \cap A_2, A_2 \setminus A_1, A_1 \setminus A_2\}$. In general there are $2^k - 1$ subsets in the slice partition, which we will denote by B_{π_j}, with n_B being the number of subsets it contains.

We now make two key observations:

1. Each observation Y_{t_j} is a sum of independent random variables $L(B_{\pi_j})$. We can see this by writing

$$Y_t = L(A_t) = L\left(\bigcup_{j:B_{\pi_j} \in \mathcal{P}(t)} B_{\pi_j}\right),$$

   where $\mathcal{P}(t)$ gives the collection of subsets in the slice partition whose union equals A_t. We can then use the additivity and independent scattering of the Lévy basis to write

$$L\left(\bigcup_{j:B_{\pi_j} \in \mathcal{P}(t)} B_{\pi_j}\right) = \sum_j L(B_{\pi_j}) \quad a.s.,$$

   where the terms $L(B_{\pi_j})$ are independent.

2. Any two observations Y_{t_i}, Y_{t_j} are conditionally independent given $L(B_{i,j})$, where $B_{i,j}$ is the union of sets in the partition giving the intersection of A_{t_i}, A_{t_j}. This follows from the independent scattering property of L. Furthermore, this property can be extended to a set of observations $Y_{t_1}, \ldots, Y_{t_n}$ by conditioning on $\{L(B_k)\}$ for the collection of sets $\{B_k\}$ giving all the intersections of $A_{t_1}, \ldots A_{t_n}$.

Note that the distribution of $L(B_{\pi_j})$ is defined via its cumulant function $|B_{\pi_j}| C(u, L')$; for many infinitely divisible distributions this corresponds to a distribution in the same class as L'.

From the above discussion we see that the observations $Y_{t_1}, \ldots, Y_{t_k}$ can be written in terms of sums of independent random variables $L(B_{\pi_j})$, which have known distributions. Thus the random variables $L(B_{\pi_j})$ constitute a tractable representation of the trawl process.

Definition 2.14: Let Y_t be a trawl process with Lévy basis L, observed at times $t_1, \ldots, t_k$, with corresponding slice partition $\{B_{\pi_j}\}$. The random variables $\{L(B_{\pi_j})\}$ are then called *slices*, and B_{π_j} are the *slice sets*.

The above discussion is summarised in the following proposition:

Proposition 2.15: A trawl process can be simulated exactly by using the slice partition, conditional on the following assumptions:

1. The area $|B_{\pi_j}|$ of each slice is known.
2. We can simulate from the distribution corresponding to the cumulant function $C(u,\cdot) = \alpha C(u, L')$ for any $\alpha > 0$.

The slice representation can be used to obtain a simplified version of the joint characteristic function for the trawl process.

Lemma 2.16: The joint characteristic function of $(Y_{t_1}, \ldots, Y_{t_k})$ is given by

$$E\left[\exp\left\{i\sum_{m=1}^{k} u_m Y_{t_m}\right\}\right] = \prod_{j=1}^{n_B} \exp\left\{\lambda(B_{\pi_j})\, C\left(\sum_{m:A_m \supset B_{\pi_j}} u_m, L'\right)\right\}.$$

Proof. By construction of the slices, we have

$$\begin{aligned} E\left[\exp\left\{i\sum_{m=1}^{k} u_m Y_{t_m}\right\}\right] &= E\left[\exp\left\{i\sum_{m=1}^{k} u_m \left(\sum_{j:B_{\pi_j}\subset A_m} L(B_{\pi_j})\right)\right\}\right] \\ &= E\left[\exp\left\{i\sum_{j=1}^{n_B} \left(\sum_{m:A_m \supset B_{\pi_j}} u_m\right) L(B_{\pi_j})\right\}\right], \end{aligned}$$

and the result then follows from the independence of $L(B_{\pi_j})$ and relation (3). □

Although the slice partition is very useful in theory, the fact that it has a size of order 2^k for k observations means that it quickly becomes computationally intractable for simulation and estimation purposes. This issue motivates the next subsection, which shows that the slice partition simplifies considerably for a monotonic trawl.

2.3.1 Monotonic trawl

We now consider the special case where the trawl set $A = A_0$ lies in $\mathbb{R}^2$ and is given by the area between the x-axis and a monotonically increasing function f on $(-\infty, 0]$, so we have

$$A_t = \{(s,\xi) : s \le t,\ \xi \le f(s-t)\}, \tag{6}$$

which is called a monotonic trawl. As will be discussed in Section 2.4, this form of trawl set is quite natural when considering causal systems.

When using a monotonic trawl, the expression for the intersection of multiple trawl sets can be simplified, as shown in the following lemma:

Lemma 2.17: If the trawl is monotonic, as specified in (6), the full collection of slice sets for time points $t_1, \ldots, t_k$ is given by

$$\left\{A^c_{t_{i-1}} \cap (A_{t_i} \cap A_{t_j}) \cap A^c_{t_{j+1}},\ 1 \le i \le j \le k,\right\} := \{S(i,j)\}, \tag{7}$$

with the convention $A_{t_0} = A_{t_{k+1}} = \emptyset$, giving a partition of size $k(k+1)/2$.

Proof. We start by considering a slice given by the intersection of trawl sets; note that this is just the set bounded on the right by the first time point t_i, and bounded above by the smallest

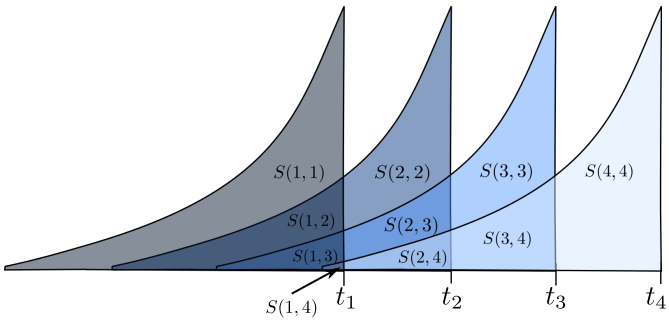


Figure 4: Illustration of slice partition.

value of $f(s-t_j)$, which occurs at the last time point t_j, so the intermediate time points are irrelevant. Specifically, for $t_i < t_j$ we have $f(s-t_j) \leq \ldots \leq f(s-t_i)$, which gives

$$A_{t_i} \cap A_{t_{i+1}} \cap \ldots \cap A_{t_j} = \{(s,\xi) : s \leq (-\infty, t_i)\ \xi \leq f(s-t_j)\} = A_{t_i} \cap A_{t_j}.$$

In particular, this implies that there are no sets in the partition of form $A_{t_i} \cap A^c_{t_k} \cap A_{t_j}$.

Another simplification occurs when considering intersections involving trawl set complements, by using $f(s-t_j) > \ldots > f(s-t_k)$ to give

$$A_{t_i} \cap A^c_{t_{i+1}} \cap \ldots \cap A^c_{t_k} = \{(s,\xi) : s \leq t_i,\ \xi \leq f(s-t_i),\ \xi > f(s-t_j)\} = A_{t_i} \cap A^c_{t_j}.$$

Similarly we obtain

$$A^c_{t_i} \cap A^c_{t_{i+1}} \cap \ldots \cap A_{t_j} = A^c_{t_{j-1}} \cap A_{t_j},$$

and the result then follows. □

We can obtain explicit expressions for the slice sets in terms of the trawl function f that bounds the trawl set above, using the definition in (7) we have

$$\begin{aligned} S(i,j) &= \{\xi, s : f(s-t_{j+1}) \leq \xi \leq f(s-t_j),\ t_{i-1} \leq s \leq t_i\}, \quad i,j \neq 1, \\ S(1,j) &= \{\xi, s : f(s-t_{j+1}) \leq \xi \leq f(s-t_j),\ -\infty \leq s \leq 0\}, \end{aligned} \tag{8}$$

as illustrated in Figure 4.

The area of these sets is given by

$$|S(i,j)| = F(t_i - t_j) - F(t_i - t_{j+1}) - F(t_{i-1} - t_j) + F(t_{i-1} - t_{j+1}), \tag{9}$$

for $i,j \neq 1$, where

$$F(t) = \int_{-\infty}^{t} f(s)\, ds.$$

2.4 Temporal dependence structure of trawl processes

We now discuss in detail how the trawl set determines the autocorrelation structure of the trawl process, which provides a basis for choosing which trawl set to use in a given model.

In the trawl process definition 2.9, we only required the trawl set A to be in $\mathcal{B}_b(\mathbb{R}^d \times \mathbb{R})$, with points in A denoted by (ξ, s). It will be useful in the following to consider the d-dimensional cross-sections of A, defined as

$$E_s = \{\xi \in \mathbb{R}^d \,\big|\, (\xi, s) \in A\}. \tag{10}$$

These cross-sections can be used to rewrite the cumulant function of the trawl process Y_t, which is given in (3) in terms of the Lévy seed L'. In particular,

$$\begin{aligned} C\left(u, L(A_t)\right) = \lambda(A_t) C\left(u, L'\right) &= \left(\int_{s:(\xi,s)\in A_t} \lambda^d(E_s)\, ds \right) C\left(u, L'\right) \\ &= \int_{s:(\xi,s)\in A_t} \lambda^d(E_s) C\left(u, L'\right)\, ds, \end{aligned} \tag{11}$$

where λ^d denotes the d-dimensional Lebesgue-measure. This representation shows that, in terms of its cumulants, $L(A)$ can be interpreted as a superposition of variables L_s, which have cumulant function

$$C\left(u, L_s\right) = \lambda^d(E_s) C\left(u, L'\right). \tag{12}$$

In this interpretation, the variable L_s is seen as the contribution from some abstract event at time s to the value of the trawl process $Y_t = L(A_t)$, and Y_t can in turn be seen as an accumulation of random pulses corresponding to L_s. This shows why it is common to consider trawl sets A_t where the last coordinate is bounded on the right by t, as this means that the value of Y_t is only influenced by events at times $s < t$, corresponding to a causal system.

For most causal systems it is reasonable to expect that the contribution of recent events will be greater than that of events far in the past, which can be stated concretely as

$$E(L_s) \leq E(L_v), \quad s < v \leq t. \tag{13}$$

Using equation (12) we get that $E(L_s) = \lambda^d(E_s)E(L')$, so the above assumption implies that the cross-section areas $\lambda^d(E_s)$ increase monotonically with s. In particular, this motivates the assumption of a monotonic trawl, as discussed in subsection 2.3.1.

2.4.1 Autocovariance function

By characterising the trawl set in terms of cross-sections we obtain a simple expression for the autocovariance function of the trawl process Y_t. Using lemma 2.12, we can differentiate with respect to u_j to show that

$$\begin{aligned} E[\Lambda_i] &= \lambda(A_{t_i})E(L'), \\ E[\Lambda_i\Lambda_j] &= \lambda(A_{t_i} \cap A_{t_j})\mathrm{Var}(L') + \lambda^2(A)E^2(L'). \end{aligned}$$

Hence the autocovariance is given by $\varphi(h) = \lambda(A_0 \cap A_h)\,\mathrm{Var}(L')$. The area of the trawl set intersection can be written in terms of the cross-sections, specifically

$$\lambda(A_0 \cap A_h) = \int_{s:(\xi,s)\in A_0\cap A_h} \lambda(E_s \cap E_{s-h})\, ds$$

This expression can be simplified if we refine the assumptions about the cross-sections mentioned above. First, we assume that the trawl set A_0 is bounded on the right, so that it only contains points (ξ, s) for $s \in (-\infty, 0]$, corresponding to a causal system. Second, we assume that the cross-sections are nested, such that $E_s \cap E_{s-h} = E_{s-h}$. This is a slightly stronger assumption than (13), which implies that the cross-sections have increasing areas, as we are now concerned with geometric properties. However, it is clear from (11) that the distribution of the trawl process Y_t is fully specified in terms of the areas of the cross-sections, and so there is no loss of generality for the trawl process in making the nesting assumption.

Having made these assumptions, we can now write

$$\begin{aligned}\varphi(h) &= \mathrm{Var}(L') \int_{-\infty}^{0} \lambda(E_{s-h})\, ds = \mathrm{Var}(L') \int_{h}^{\infty} \lambda(E_{-v})\, dv, \\ &= \mathrm{Var}(L') \left(\lambda(A) - \int_{0}^{h} \lambda(E_{-v})\, dv \right).\end{aligned}$$

This expression can be used to solve for the area $\lambda(E_v)$ in terms of the autocovariance function by differentiating the above expression to get

$$\lambda(E_{-h}) = -\varphi'(h)/\mathrm{Var}(L'), \quad h \geq 0.$$

Since the area $\lambda(E_{-h})$ must be non-negative, and by assumption is monotonically decreasing with increasing h, it follows that $\varphi'(\cdot)$ must be non-positive and have a magnitude that is monotonically decreasing.

The above relation shows that given a monotonically decreasing autocovariance function $\varphi(\cdot)$ with an absolutely decreasing derivative, we can construct E_s and the corresponding trawl set A such that the process $Y_t = L(A_t)$ has autocovariance function φ, as shown in Barndorff-Nielsen et al. (2014) for integer-valued trawl processes with trawl sets $A \in \mathbb{R} \times \mathbb{R}$.

3 Latent trawl process model

In this section we use the framework of trawl processes to construct a new model for extreme values observed over time, which extends the hierarchical setup of Bortot and Gaetan (2014) in several directions.

The starting point for the model is the observation that the marginal distribution of exceedances converges to the generalised Pareto distribution (GPD), see Davison and Smith (1990). In Bortot and Gaetan (2014), a decomposition of the GPD is used to construct a hierarchical model for exceedances that preserves this distribution marginally. We adopt this hierarchical structure, and then proceed to introduce a new and more flexible model by using the properties of trawl processes described in Section 2.

The original approach of Bortot and Gaetan (2014) involved using a Markov chain to generate dependence in the exceedances; by using a trawl process instead of a Markov chain we obtain a model with a more flexible dependence structure. Moreover, the trawl process framework provides a unified procedure for generating processes with a given infinitely divisible distribution and autocovariance function, whereas Bortot and Gaetan (2014) consider two different specifications for the latent Markov chain with the same marginal distribution and autocovariance function.

We also consider a modification of the new model that enables it to be used for any generalised Pareto distribution, thus removing the restriction $\xi > 0$ on the shape parameter that was inherited from the original model of Bortot and Gaetan (2014). This means that the modified model can be used for processes with less heavy-tails, which are often found in environmental applications, as illustrated by the application to pollution data in Section 4.4. The modification also improves the interpretability of the parameters, and appears to make the estimation procedure more efficient.

3.1 Basic structure

This subsection introduces the hierarchical structure used for the extreme value model, which is taken from Bortot and Gaetan (2014). Consider a time-series $\{Y_j\}_{j=1,\ldots,k}$, assumed to be

strictly stationary. Our aim is to model the observations where the value of Y_j is considered to be extreme, meaning that $Y_j > u$ for a fixed threshold u.

In order to focus on the extreme values, we will only consider the values and occurrence times of exceedances. To this end, define the exceedances X_j by

$$X_j := \max(Y_j - u, 0), \quad j = 1, \ldots, k.$$

From standard extreme value theory (Pickands, 1975; Davison and Smith, 1990), assuming $\{Y_j\}$ are in the domain of attraction of some extreme value distribution, the conditional exceedances $\{X_j | X_j > 0\}$ converge to a generalised Pareto distribution (GPD) as $u \to \infty$. Based on this result we will assume that the conditional distribution of X_j given $X_j > 0$ can be approximated by a GPD, for a sufficiently large threshold u. The density of the GPD is written as

$$f_{\text{GPD}}(x|\alpha, \beta) = \frac{\alpha}{\beta}\left(1 + \frac{x}{\beta}\right)_+^{-(\alpha+1)} \qquad x \geq 0,\ \alpha, \beta > 0,$$

which is a reparameterisation of the standard density with shape parameter $\xi = 1/\alpha$ and scale parameter $\sigma = \beta/\alpha$.

Following Bortot and Gaetan (2014) we now observe that the GPD can be represented as a mixture of exponential random variables with Gamma distributed parameters, motivating a hierarchical specification for the exceedance process $\{X_j\}$. In particular, the distribution of an observation X_j depends on the value of a latent process Λ at time t_j, denoted Λ_j. This latent process determines both the probability of observing an exceedance, corresponding to $X_j > 0$, and the distribution of the exceedances.

Specifically, we assume that the density of X_j is

$$f(x_j \,|\lambda_j) = \begin{cases} 1 - e^{-\kappa\lambda_j}, & x_j = 0, \\ e^{-\kappa\lambda_j}\lambda_j e^{-\lambda_j x_j}, & x_j > 0. \end{cases} \tag{14}$$

where the density f is defined with respect to a measure μ such that $\mu(dx) = \delta_0(dx) + dx$. This construction shows that conditionally on the value of Λ_j, the exceedance X_j is generated by a two-stage process: at the first stage X_j is set to zero with probability $1 - e^{-\kappa\lambda_j}$; at the second stage the distribution of X_j given $\{X_j > 0, \Lambda_j = \lambda_j\}$ is exponential with parameter λ_j. The latent process Λ may be interpreted as an inverse intensity, as higher values of Λ give a lower probability of exceeding the threshold u and a smaller expected value of exceedances.

We also require that the observations X_j are independent for distinct j, conditional on the corresponding values of Λ. Hence the conditional joint distribution of $(X_1, \ldots, X_k)$ can be written as a product:

$$f(x_1, \ldots, x_k \,|\lambda_1, \ldots, \lambda_k) = \prod_{j=1}^{k} f(x_j|\lambda_j). \tag{15}$$

This specification implies that any dependence between observations $X_1, \ldots, X_k$ comes from the dependence between corresponding elements of the latent process Λ.

To complete the GPD mixture construction, the latent process is required to have a Gamma marginal law, i.e.

$$\begin{aligned} &\Lambda_j \sim \text{Gamma}(\alpha, \beta), \quad \alpha, \beta > 0, \\ &C(u, \Lambda_j) = -\alpha \log(1 - iu/\beta), \end{aligned} \tag{16}$$

where C is the cumulant function as defined in Proposition 2.11. This specification introduces the restriction $\alpha > 0$, which means that the model cannot capture data belonging to the Weibull distribution class, corresponding to a negative shape parameter of the GPD. Section 3.4 presents a modified version of the model that removes this restriction.

It remains to show that when using the above specification, the exceedances $\{X_j : X_j > 0\}$ have a GPD marginal law. This is done by integrating out the value of Λ to get the marginal distribution of X, where we drop subscripts in the following calculations to ease notation. Specifically, consider a set $B \in \mathcal{B}(\mathbb{R})$, then we have

$$P(\{X \in B, X > 0\}) = \int_B \int_0^\infty f(x, \lambda)\, d\lambda\, dx = \int_B \int_0^\infty \lambda e^{-\lambda(\kappa+x)} f(\lambda)\, d\lambda\, dx.$$

This can be written as an expectation involving Λ to give

$$\begin{aligned} P(\{X \in B \cap X > 0\}) &= \int_B E_\Lambda[\Lambda \exp\{-\Lambda(x+\kappa)\}]\, dx \\ &= \int_B -\frac{\partial}{\partial u} E_\Lambda[\exp\{-\Lambda u\}]\Big|_{u=x+\kappa} = \int_B \frac{\alpha}{\beta}\left(1 + \frac{x+\kappa}{\beta}\right)^{-(\alpha+1)} dx, \end{aligned}$$

where the cumulant function of Λ was used in the last step. Similarly, the unconditional probability of observing an exceedance is given by

$$P(X > 0) = E_\Lambda[e^{-\kappa\Lambda}] = \left(1 + \frac{\kappa}{\beta}\right)^{-\alpha}, \tag{17}$$

and hence the distribution of the exceedances is given by

$$P(X \in B | X > 0) = \int_B \frac{\alpha}{\beta + \kappa}\left(1 + \frac{x}{\beta + \kappa}\right)^{-(\alpha+1)},$$

showing that the conditional distribution of $\{X\,|X > 0\}$ is given by $\text{GPD}(\alpha, \beta + \kappa)$.

3.2 Latent trawl process

The previous subsection describes a general hierarchical model setup, using the same structure as in Bortot and Gaetan (2014). So far, the latent process Λ has only been specified as having a Gamma marginal law. Now we depart from the approach used in Bortot and Gaetan (2014), where the latent process is assumed to be a Markov chain; rather, we consider a new model where Λ is a trawl process, as introduced in Section 2.

The conditional independence assumption of the hierarchical model means that any dependence between observations comes from the latent process, and hence this process should have a flexible dependence structure. This motivates the use of a trawl process, which can capture a wide range of dependence structures, as discussed in subsection 2.4. Using a trawl process also means that the observations X_j can be seen as coming from a continuous-time process (X_t), which is useful for statistical applications where there may be missing or irregularly spaced data.

Bortot and Gaetan (2014) consider two particular classes for the latent Markov chain, the Gaver and Lewis process (G-LP) and the Warren process (WP), and proceed to show that these two classes result in different asymptotic properties of the extremes even though they have the same autocorrelation function. In contrast, the latent trawl process in our model is specified by its trawl set, which corresponds to a particular autocorrelation function; as will be shown in Section 4.2 this process can be classified as asymptotically independent, where the form of the asymptotic dependence structure is influenced by the trawl set.

3.2.1 Marginal distribution

Section 2 introduced the definition of a Lévy basis and discussed how to construct a particular Lévy basis L from a random variable L', called the Lévy seed. We now use this approach to obtain a trawl process Λ with a Gamma marginal law, following Proposition 2.11. Specifically, we fix a trawl set A in $\mathbb{R}^d \times \mathbb{R}$, and let the Lévy seed have a normalised Gamma distribution, i.e.

$$L' \sim \text{Gamma}(\alpha/|A|, \beta), \tag{18}$$

then the trawl process defined by $\Lambda_t = L(A_t)$ has a Gamma(α, β) distribution.

Combining the trawl process Λ with the hierarchical model presented in the previous subsection, we obtain a stochastic process (X_t) with finite-dimensional densities given by

$$f(x_1, \ldots, x_k) = \int_{\mathbb{R}_+^k} \prod_{j=1}^{k} \left\{ \begin{array}{ll} 1 - e^{-\kappa\lambda_j}, & x_j = 0 \\ \lambda_j e^{-\lambda_j(\kappa + x_j)}, & x_j > 0 \end{array} \right\} dF(\lambda_1, \ldots, \lambda_k), \tag{19}$$

using the independence assumption in (15). These densities depend on the joint distribution of $(\Lambda_1, \ldots, \Lambda_k)$, which is fully specified by the trawl set A and the Lévy seed L', as shown in lemma 2.16.

3.2.2 Trawl set

To complete the definition of the trawl process Λ it now remains to specify the trawl set A. For our model we use the so-called exponential trawl set; this is the set in $\mathbb{R} \times (-\infty, 0]$ that is bounded below by the x-axis and above by the function $f(x) = \exp(\rho x)$ for some $\rho > 0$. The resulting process is then called the *exponential trawl process*.

As discussed in Section 2.4, the cross-section E_x of the trawl set can be associated with a random variable L_x, which is interpreted as the contribution to the value of Λ_0 coming from events x time units in the past. The expected value of L_x is scaled by the Lebesgue measure of the cross-section, seen as a subset of $\mathbb{R}^d$. For the exponential trawl, the cross-sections are vertical intervals, so

$$E_x = \{(x, y) : y \leq e^{\rho x}\}, \quad x \leq 0,$$

with one-dimensional Lebesgue-measure $e^{\rho x}$.

The autocovariance function of the trawl process is given by $\varphi(h) = \lambda(A \cap A_h)\text{Var}(L')$, and using the exponential trawl gives

$$\lambda(A) = \int_{-\infty}^{0} f(s)ds = \frac{1}{\rho}, \quad \lambda(A \cap A_h) = \int_{-\infty}^{0} f(s-h)ds = \frac{e^{-\rho h}}{\rho}.$$

Combining with (18) gives $\text{Var}(L') = (\alpha\rho)/\beta^2$, resulting in the autocovariance function

$$\varphi(h) = e^{-\rho h}\frac{\alpha}{\beta^2} = e^{-\rho h}\text{Var}(\Lambda_t).$$

Thus the autocorrelation function of the exponential trawl process has the same shape as the trawl set, which is a particular property of the exponential trawl.

We can also consider a general exponential trawl set, which is constructed from linear combinations of basic exponential trawls. Specifically, define the general exponential trawl set of order p to be bounded above by the function

$$f_p(x) = \sum_{i=1}^{p} w_i e^{\rho_i x},$$

with $\sum_i w_i = 1$. The latter restriction is necessary to make the parameters identifiable, as any scaling factor in the weights w_i will scale the area of the trawl, and thus be cancelled by the normalisation in (18). When using the general exponential trawl set, the resulting trawl process has an autocorrelation function $r(h)$ given by the corresponding linear combination of $e^{-\rho_i h}$. This construction can be seen as a special case of a superposition-type trawl where the decay parameter ρ is randomised, as in Barndorff-Nielsen et al. (2014, Section 4); see also Barndorff-Nielsen (2001) for a similar approach applied to Ornstein-Uhlenbeck-type processes.

Summarising the above discussion, we have constructed a trawl process Λ with a marginal Gamma distribution and an autocorrelation function based on exponential decays. Using this as the latent process in the hierarchical structure results in a continuous-time process (X_t), where the exceedances $\{X_t > 0\}$ have a generalised Pareto distribution. Furthermore, the model allows for dependence between observations $\{X_t\}$, which is derived directly from the dependence in the latent trawl process.

3.3 Autocovariance structure

We now consider the autocovariance structure of the exceedance process (X_t) from the latent model. As X is a stationary process, it has autocovariance function

$$\varphi(h) = E[X_0 X_h] - E^2[X].$$

The expected value of X is obtained by conditioning on the events $\{X = 0\}$ and $\{X > 0\}$, with probabilities given by (17), and using that $\{X|X > 0\}$ has a GPD$(\alpha, \beta + \kappa)$ distribution, which gives

$$E[X] = (1 + \kappa/\beta)^{-\alpha}(\beta + \kappa)/(\alpha - 1), \quad \alpha > 1.$$

Conditioning similarly for the variance gives

$$\operatorname{var}(X) = p_1 \frac{(\beta + \kappa)^2 \left[2(\alpha - 1) - p_1(\alpha - 2)\right]}{(\alpha - 1)^2(\alpha - 2)}, \quad \alpha > 2,$$

with $p_1 = P(X > 0) = (1 + \kappa/\beta)^{-\alpha}$. The expectation of the product is more complicated, as it involves the joint density of the corresponding values of the latent trawl process:

$$\begin{aligned} E[X_0 X_h] &= 0 + \int_{\mathbb{R}_+^2} x_0 x_h \left(\int_{\mathbb{R}_+^2} f(x_0, x_h|\lambda_0, \lambda_h) f(\lambda_0, \lambda_h)\, d\lambda_0 d\lambda_h \right) dx_0\, dx_h \\ &= \int_{\mathbb{R}_+^2} x_0 x_h \left(\int_{\mathbb{R}_+^2} \lambda_0 e^{-\lambda_0(\kappa + x_0)} \lambda_h e^{-\lambda_h(\kappa + x_h)} f(\lambda_0, \lambda_h)\, d\lambda_0\, d\lambda_h \right) dx_0\, dx_h. \end{aligned}$$

Using Fubini's theorem to repeatedly exchange the order of integration gives

$$\begin{aligned} E[X_0 X_h] &= \int_{\mathbb{R}_+^2} \lambda_0 \lambda_h e^{-(\lambda_0 \kappa + \lambda_h \kappa)} f(\lambda_0, \lambda_h) \left(\int_{\mathbb{R}_+^2} x_0 x_h e^{-(x_0 \lambda_0 + x_h \lambda_h)}\, dx_0\, dx_h \right) d\lambda_0\, d\lambda_h \\ &= \int_{\mathbb{R}_+^2} \lambda_0 \lambda_h e^{-(\lambda_0 \kappa + \lambda_h \kappa)} f(\lambda_0, \lambda_h) \left(\frac{1}{\lambda_0^2 \lambda_h^2} \right) d\lambda_0\, d\lambda_h \\ &= E_{\Lambda_0, \Lambda_h} \left[\frac{e^{-(\Lambda_0 \kappa + \Lambda_h \kappa)}}{\Lambda_0 \Lambda_h} \right] = \int_\kappa^\infty \int_\kappa^\infty E_{\Lambda_0, \Lambda_h} \left[e^{-(\Lambda_0 u_0 + \Lambda_h u_h)} \right] du_0\, du_h. \end{aligned}$$

The joint expectation can now be found using the partitioning method described in section 2.3. Specifically, partition the trawl sets A_0, A_h into disjoint sets

$$B_0 = A_h \setminus A_h,\ B_{0,h} = A_0 \cap A_h,\ B_h = A_h \setminus A_0,$$

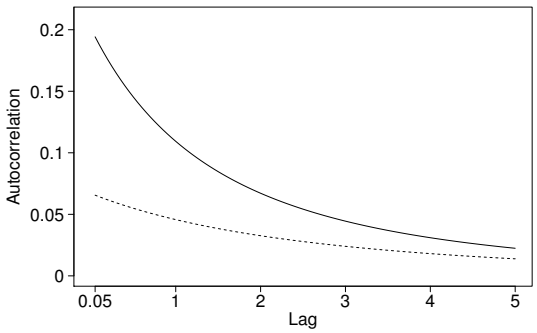


Figure 5: Autocorrelation functions of latent trawl models, the solid line corresponds to a model with $\alpha = 4, \beta = 4$, the dashed line to $\alpha = 9, \beta = 1$. Both models set $\lambda = 0.2$ and κ such that the probability of an exceedance equals 0.05.

and define the slice variables $\{S_i = L(B_i)\}$, where L is the Lévy basis derived from (18); in particular these variables are independent. Furthermore, we have $\Lambda_0 \overset{d}{=} S_0 + S_{0,h}$ and $\Lambda_h \overset{d}{=} S_{0,h} + S_h$.

Let b_i be the shape parameter of S_i, which is given by $\alpha/|A|$ times the Lebesgue measure of set B_i, then we can rewrite the above expression as

$$\begin{aligned} E[X_0 X_h] &= \int_\kappa^\infty \int_\kappa^\infty E_{S_0, S_{0,h}, S_h} \left[\exp\{-\left[u_0 S_0 + (u_0 + u_h) S_{0,h} + u_h S_h\right]\}\right] du_0 \, du_h \\ &= \int_\kappa^\infty \int_\kappa^\infty \left(1 + \frac{u_0}{\beta}\right)^{-b_0} \left(1 + \frac{u_0 + u_h}{\beta}\right)^{-b_{0,h}} \left(1 + \frac{u_h}{\beta}\right)^{-b_h} du_0 \, du_h. \end{aligned}$$

The above integral can be done numerically to obtain the autocovariance of X for given parameters $(\alpha, \beta, \rho, \kappa)$, where the areas b_i are functions of ρ and h.

The autocorrelation function given above for the exceedance process X is discontinuous at zero, which is a consequence of the mixture construction; the latent trawl process itself is continuous in probability.

As remarked in section 2.2.1, the trawl process separates the parameters controlling the marginal and dependence properties of the model. However, this is not the case when considering the full hierarchical model, as the parameters α, β and κ in the marginal distribution also influence the autocovariance structure of the process. This is illustrated in figure 5, which shows two different autocorrelation functions obtained by varying the parameters α and β. This conflation of marginal and dependence parameters motivates the model in the following subsection.

3.4 Marginal transformation model

This subsection considers a modification of the latent trawl model that has the effect of separating the marginal and dependence properties. This modification removes the restriction of the model to GP distributions with positive shape parameter, corresponding to the Fréchet class of extreme value distributions. This parameter restriction was highlighted as a potential problem in the conclusion of Bortot and Gaetan (2014), where a similar modification was suggested but not explored further.

The model resulting from this modification is also easier to interpret, as the role of each parameter is uniquely defined in terms of controlling either the marginal distribution, probability of exceedance or dependence properties. This also contributes to identifiability of the parameters, in particular we found that the estimation procedure appears to be more effective for the modified model.

The modified model is derived from the original model in two steps: the first is to fix the parameters α, β of the latent Gamma distribution such that only the parameters associated with the trawl set will influence the trawl process, and thus also the dependence of the exceedances. In the following we will work with $\alpha = \beta = 1$, such that the exceedances have marginal law given by $GPD(1, 1+\kappa)$.

The second step is to add an extra layer to the modified model, which uses a standard probability integral transform to give the marginals a $GPD(\xi, \sigma)$ distribution, specifically

$$Z_j = F^{-1}_{\mathrm{GPD}(\xi,\sigma)}(F_{\mathrm{GPD}(1,1+\kappa)}(X_j)) := g(X_j) \quad \xi \in \mathbb{R}, \sigma > 0.$$

This modified version will be called the marginal transformation (MT) model; it has finite-dimensional densities given by

$$f(z_1, \ldots, z_k) = \int_{\mathbb{R}^k_+} \prod_{j=1}^{k} \left\{ \begin{array}{ll} 1 - \exp\{-\kappa\lambda_j\}, & z_j = 0 \\ J(z_j)\, \lambda_j \exp\{-\lambda_j(\kappa + g^{-1}(z_j))\}, & z_j > 0 \end{array} \right\} dF(\lambda_1, \ldots, \lambda_k),$$

$$J(z_j) = \frac{f_{\mathrm{GPD}(\xi,\sigma)}(z_j)}{f_{\mathrm{GPD}(1,1+\kappa)}(g^{-1}(z_j))},$$

with respect to the measure $\mu(dx) = \delta_0(dx) + dx$. This follows by noting that the transformation g maps the event $\{X = 0\}$ to $\{Z = 0\}$, thus leaving the atom at zero unchanged, whereas the transformation of the continuous part on $\{X > 0\}$ introduces a standard Jacobian term.

4 Model fitting and evaluation

4.1 Pairwise likelihood

We now consider parameter estimation for the latent trawl model described above. Given observations $\{Y_j\}$ from a stationary time series, we transform them as in Section 3.1 to get exceedances $X_j := \max(Y_j - u, 0)$, $j = 1, \ldots, k$. We assume there are l positive observations $X_{p_1}, \ldots, X_{p_l}$, and m observations $X_{q_1}, \ldots, X_{q_m}$ taking the value zero, which we will call exceedances and non-exceedances, respectively.

The likelihood of the observations $\{X_j\}$ under the original latent trawl model now follows from (19), and can be written as

$$f(x_1, \ldots, x_k) = \int_{\mathbb{R}^k_+} \prod_{r=1}^{l} \lambda_{p_r} \exp\{-(x_{p_r} + \kappa)\lambda_{p_r}\} \prod_{s=1}^{m} (1 - \exp\{-\kappa\lambda_{q_s}\})\, dF(\lambda_1, \ldots, \lambda_k),$$

where F is the joint density function of the trawl process observations $\Lambda_1, \ldots, \Lambda_k$. The integrand above can be expanded to obtain a sum involving 2^m number of terms, by first defining $\mathcal{S}_t$ as the collection of subsets of $\{q_1, \ldots, q_m\}$ of size t, and letting $u_r = x_{p_r} + \kappa$, to give

$$\sum_{t=0}^{m} (-1)^t \sum_{\pi_t \in \mathcal{S}_t} \prod_{r=1}^{l} \lambda_{p_r} \exp\left\{-u_r \lambda_{p_r}\right\} \prod_{s_j \in \pi_t} \exp\left\{-\kappa\lambda_{s_j}\right\},$$

with the convention $\prod_{s_j \in \pi_0}(\ldots) = 1$. This can be rewritten as

$$= \sum_{t=0}^{m} \sum_{\pi_t \in \mathcal{S}_t} (-1)^{t+l} \frac{\partial}{\partial u_1} \ldots \frac{\partial}{\partial u_l} \exp\left\{ -\sum_{r=1}^{l} u_r \lambda_{p_r} - \sum_{s_j \in \pi_t} \kappa\lambda_{s_j} \right\}.$$

Using this expression and exchanging integrals and partial derivatives we see that the full likelihood reduces to

$$\sum_{t=0}^{m} \sum_{\pi_t \in \mathcal{S}_t} (-1)^{t+l} \frac{\partial}{\partial u_1} \dots \frac{\partial}{\partial u_l} \int_{\mathbb{R}_+^k} \exp\left\{ -\sum_{r=1}^{l} u_r \lambda_{p_r} - \sum_{s_j \in \pi_t} \kappa \lambda_{s_j} \right\} dF(\lambda_1, \dots, \lambda_k)$$
$$= \sum_{t=0}^{m} \sum_{\pi_t \in \mathcal{S}_t} (-1)^{t+l} \frac{\partial}{\partial u_1} \dots \frac{\partial}{\partial u_l} E\left[\exp\left\{ -\sum_{r=1}^{t+l} u_r \lambda_{p_r} - \sum_{s_j \in \pi_t} \kappa \lambda_{s_j} \right\} \right],$$

where the expectation is with respect to the corresponding variables $\Lambda_{p_1}, \dots, \Lambda_{p_l}, \Lambda_{s_1}, \dots, \Lambda_{s_t}$ determined by π_t.

The expected values in these terms are just joint characteristic functions, as given in Lemma 2.16, which involve the parameters of the Lévy seed L and trawl intersection areas, hence the complete likelihood reduces to a sum of partial derivatives of characteristic functions.

Although the likelihood is available analytically, it is not easy to compute in pratice, for two reasons. First, it would require multiple numerical partial derivatives to be performed. Second, the number of non-exceedances m is usually close to the number of observations k; this is because the latent model is defined to have a GP distribution, and to justify this assumption we need to consider a sufficiently high threshold for exceedances, often given by a large percentile of the observations. Now the two first sums in the likelihood above have 2^m terms in total, and hence the likelihood becomes computationally intractable for any reasonable sample size k.

Because of the computational issues associated with the sample size, we consider using a pairwise likelihood approach for model fitting, which is a particular kind of composite likelihood (Varin, 2008; Cox and Reid, 2004; Varin et al., 2011). As stated in Varin (2008), composite likelihood estimators have good properties when the data can be seen as consisting of roughly independent blocks, i.e. the autocorrelation function decays sufficiently fast. Thus using the pairwise likelihood should provide reasonable performance for the latent trawl model with an exponential trawl set.

Given observations $x_1, \dots, x_k$, the pairwise likelihood f_{PL} takes the form

$$f_{PL}^{\triangle}(x_1, \dots, x_k) = \prod_{i=1}^{k-1} \prod_{j=i+1}^{\min(i+\triangle, k)} f(x_i, x_j),$$

where $f(\cdot)$ is the original likelihood function and $\triangle$ denotes the maximum separation between observations. For the latent trawl model, each pairwise likelihood term $f(x_i, x_j)$ involves at most 4 terms in the sum above, and so these terms can be evaluated explicitly in terms of the parameters of the trawl process. There are four different cases, as each of x_i and x_j can be an exceedance or not; the explicit forms of $f(x_i, x_j)$ are given in the Appendix.

4.2 Dependence measure φ

We now consider the extremal dependence structure of our model. A common measure of dependence at high levels is the extremal index (Leadbetter et al., 1983), which can be characterised as

$$\theta = \lim_{n \to \infty} P(X_i \leq u_n, 2 \leq i \leq l_n | X_1 > u_n),$$

where u_n is an increasing sequence of thresholds and $l_n = o(n)$ (Ancona-Navarrete and Tawn (2000), following O'Brien (1987)).

The extremal index θ essentially describes the dependence across blocks of observations whose length tends to infinity; it can also be defined as the reciprocal cluster length, where a cluster is

a collection of exceedances in a block. Thus to estimate the extremal index one has to consider very high-dimensional joint distributions, which makes it analytically intractable in many cases, in particular for our latent trawl model. For the applications in the following subsections we will instead consider simulation-based estimates of the extremal index.

There are also measures of extremal dependence that work on a shorter range of observations than the extremal index. Coles et al. (1999) quantify the dependence between the extreme values of two random variables X_1, X_2 in the upper tail-dependence coefficient, given by

$$\chi = \lim_{u \to 1} P(F(X_2) > u | F(X_1) > u),$$

where F is the common marginal distribution of X_1, X_2. In other words, χ gives the limiting probability of X_2 exceeding the threshold u given an exceedance X_1, with both variables on a uniform scale. When X_1, X_2 come from a stationary time series, χ can be seen as the probability of observing consecutive exceedances given a single exceedance. To get a broader characterisation of the extremal dependence one can also consider the complete function

$$\chi(u_1, u_2) = P(F(X_2) > u_2 | F(X_1) > u_1),$$

defined on $[0,1]^2$, , where $\chi(u_1, 0) = 1$ and $\chi(0, u_2) = 1 - u_2$.

We would like to use the function $\chi(u_1, u_2)$ and the limiting measure χ to evaluate the dependence between two observations X_1, X_2 from the latent trawl model as a function of the lag $t_2 - t_1$. However, the above definition is based on the assumption that X_1, X_2 are continuously-valued random variables with distribution function F defined on their entire range, and this does not fit with the exceedance framework where the limiting distribution can only be assumed to hold above a threshold. Some extreme value models (see e.g. Ledford and Tawn (1996)) do not consider exceedances separately, but specify the same probability density function f for the whole range of observations, and then treat observations below the threshold u as censored, so they have probability

$$\int_0^u f(s) ds.$$

Our model differs in that it models the occurence of exceedances explicitly, resulting in a distribution for exceedances X only with an atom at zero, and hence the standard definition of χ cannot be applied directly. However, we can construct an analogue to the tail dependence coefficient by conditioning on both exceedances being positive, but care must be taken to ensure that the resulting measure is uniform on the u_1 margin, as shown in the following definition.

Definition 4.1: For the latent model as defined in Section 3.1, we define the *conditional tail dependence coefficient* as

$$\varphi = \lim_{u_1, u_2 \to 1} \varphi(u_1, u_2) := \lim_{u_1, u_2 \to 1} P(F_{2e}(X_2) > u_2 \,|F_{2e}(X_1) > u_1, X_1 > 0, X_2 > 0), \tag{20}$$

where

$$F_{2e}(x) = P(X_1 \leq x \,|X_1 > 0, X_2 > 0) = P(X_2 \leq x \,|X_1 > 0, X_2 > 0).$$

The conditional tail dependence function $\varphi(u_1, u_2)$ can be calculated explicitly in terms of the parameters of the latent trawl model. Specifically, by conditioning on the latent trawl process we obtain

$$\begin{aligned}
&P(F_{2e}(X_1) > u_1, F_{2e}(X_2) > u_2, X_1 > 0, X_2 > 0) \\
=&P(X_1 > F_{2e}^{-1}(u_1), X_2 > F_{2e}^{-1}(u_2), X_1 > 0, X_2 > 0) \\
=&\int \exp\left\{-\lambda_1 (F_{2e}^{-1}(u_1) + \kappa)\right\} \exp\left\{-\lambda_2 (F_{2e}^{-1}(u_2) + \kappa)\right\} f(\lambda_1, \lambda_2)\, d\lambda_1 d\lambda_2 \\
=&E_{\Lambda_1, \Lambda_2} \exp\left\{-\Lambda_1 (F_{2e}^{-1}(u_1) + \kappa) - \Lambda_2 (F_{2e}^{-1}(u_2) + \kappa)\right\}.
\end{aligned}$$

This expectation can be calculated by using the partition representation of the trawl process, similar to the calculation in section 3.3, to give

$$\begin{aligned}&P(F_{2e}(X_1) > u_1, F_{2e}(X_2) > u_2, X_1 > 0, X_2 > 0)\\ &= \left(1 + \frac{\kappa + F_{2e}^{-1}(u_1)}{\beta}\right)^{-b_1} \left(1 + \frac{2\kappa + F_{2e}^{-1}(u_1) + F_{2e}^{-1}(u_2)}{\beta}\right)^{-b_2} \left(1 + \frac{\kappa + F_{2e}^{-1}(u_2)}{\beta}\right)^{-b_3}.\end{aligned}$$

Setting $u_2 = 0$ in this expression and noting that $F_{2e}^{-1}(0) = 0$ shows that

$$\begin{aligned}&P(F_{2e}(X_1) > u_1, X_1 > 0, X_2 > 0)\\ &= \left(1 + \frac{\kappa + F_{2e}^{-1}(u_1)}{\beta}\right)^{-b_1} \left(1 + \frac{2\kappa + F_{2e}^{-1}(u_1)}{\beta}\right)^{-b_2} \left(1 + \frac{\kappa}{\beta}\right)^{-b_3},\end{aligned} \tag{21}$$

giving the denominator of φ, so that we get

$$\varphi(u_1, u_2) = \left(1 + \frac{F_{2e}^{-1}(u_2)}{\beta + 2\kappa + F_{2e}^{-1}(u_1)}\right)^{-b2} \left(1 + \frac{F_{2e}^{-1}(u_2)}{\beta + \kappa}\right)^{-b3}.$$

We can also write

$$F_{2e}(x) = 1 - \frac{P(X_1 > x, X_1 > 0, X_2 > 0)}{P(X_1 > 0, X_2 > 0)},$$

where the probabilities can be derived from (21) by the substitution $u_1 \to F_{2e}(x)$, resulting in

$$F_{2e}(x) = 1 - \left(1 + \frac{x}{\beta + 2\kappa}\right)^{-b2} \left(1 + \frac{x}{\beta + \kappa}\right)^{-b3}.$$

In particular, we now get the same marginal scaling as the original χ, namely $\varphi(u_1, 0) = 1, \varphi(0, u_2) = 1 - u_2$.

Using this expression, we see that for the original latent trawl model, the limiting coefficient as $u_1, u_2 \to 1$ equals zero, so according to this measure the model is asymptotially independent. Furthermore, the decay to zero increases with the value of b_3; in other words, the larger the intersection of the trawl sets given by X_1, X_2, the slower the model decays to independence as the threshold increases.

4.3 Heathrow data

In this subsection we use the latent trawl model to analyse a data set consisting of hourly accumulated rainfall amounts at Heathrow (UK) over the years 1980-2012, provided by UK Meteorological Office (2012). We set the threshold u at the 95 percentile of the original data (8.9mm), resulting in the time series of exceedance values shown in Figure 6.

We fitted both the latent trawl model described in Section 3.2 and the latent Markov chain model of Bortot and Gaetan (2014), using both the G-LP and Warren process for the Markov chain. The parameter estimation was done using pairwise likelihood as described in Section 4.1, with separation parameter $\delta = 4$; the resulting estimates are shown in Table 1. We see that the marginal parameters are similar across the models, which is reasonable as the models all have a marginal GPD$(\alpha, \beta + \kappa)$ distribution and κ controls the marginal exceedance probability.

The parameter ρ controls the latent dependence structure of the models. For the latent trawl process it is the decay parameter of the exponential trawl function, whereas for the latent Markov chains it enters in the autocorrelation function $\phi(h) = \rho^h = \exp(h \log(\rho))$. Using this relation we see from the fitted values of ρ that all three latent processes have similar autocorrelation functions, where the G-LP and WP processes have the fastest and slowest decay, respectively.

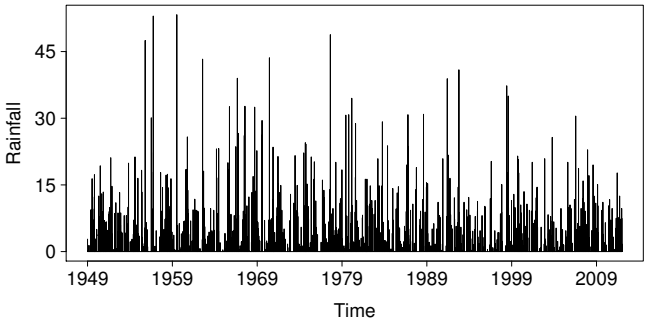


Figure 6: Heathrow rainfall exceedances.

	α	β	ρ	κ
Latent trawl	6.33	20.12	0.27	12.18
G-LP	6.43	20.64	0.70	12.25
WP	6.30	19.94	0.78	12.15

Table 1: Estimated parameters for Heathrow rainfall data using latent trawl model and latent Markov chain model with G-LP and WP chains.

Figure 7 shows estimates of the extremal index for the latent trawl and latent Markov chain models. These estimates are based on simulating time-series of length 1,000,000 from the fitted models, and then estimating θ as the inverse cluster length using the R package `evd`, where a cluster is defined as ending when three consecutive exceedances fall below the threshold.

These estimates indicate that the latent trawl and latent Markov chain models both manage to capture the main dependence structure in the extremes. The only visible difference occurs at high thresholds, where the G-LP model appears to underestimate θ, i.e. overestimate the dependence. This discrepancy at high levels could be explained by the fact that the G-LP model is asymptotically dependent, as shown in Bortot and Gaetan (2014), whereas the empirical estimates of θ indicate that the rainfall data is asymptotically independent.

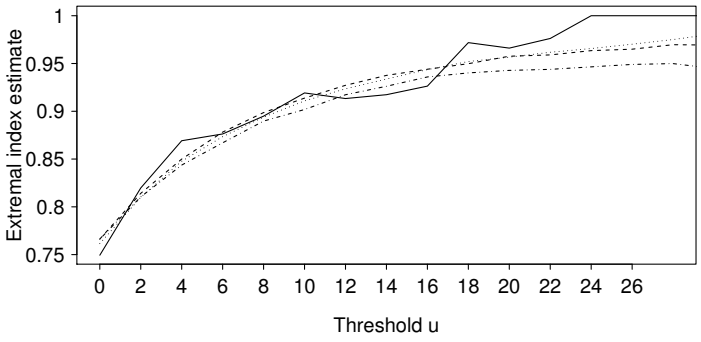


Figure 7: Estimated extremal index for empirical rainfall data (solid), latent trawl (dashed), G-LP (dash-dotted) and WP (dotted).

4.4 Pollution data

We now consider a second application that illustrates the marginal transformation model introduced in 3.4. We use this model to analyse measurements of Ozone levels in Bloomsbury, London; specifically, the data gives the daily maximum of the 8-hour running mean, measured in units of $\mu g/m^3$ (microgrammes per cubic metre). They are obtained from the Environmental

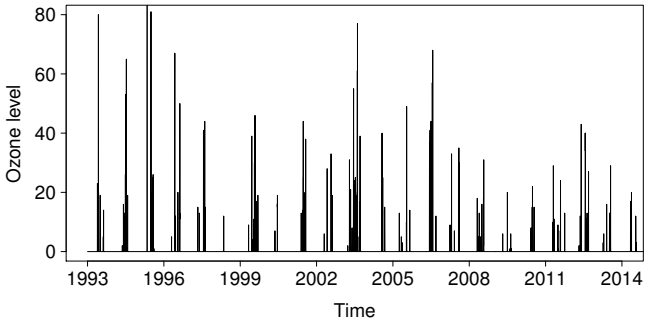


Figure 8: Ozone level exceedances

Research Group, King's College London (2015).

Although there is evident seasonality in the original data, the effect diminishes significantly as the threshold increases, and so we do not adjust for seasonality in the extreme values for this application. We chose the threshold u to be the 97 percentile of the original data ($81\mu g/m^3$), resulting in the exceedances shown in Figure 8.

Fitting the GPD distribution to the data directly indicates a negative value of the shape parameter (α or ξ), i.e. a finite upper bound for the distribution, which cannot be captured by the standard model. Thus we use the transformed model described in section 3.4, which corresponds to taking the basic hierarchical model (either based on latent trawl or latent Markov chain as in Bortot and Gaetan (2014)), fixing $\alpha, \beta = 1$, and then transforming the marginals to GPD(ξ, σ).

The models were fitted using pairwise likelihood with $\triangle = 4$. Table 2 shows the resulting estimates for the marginal transformed versions of the latent trawl and latent Markov chain models.

Figure 9 shows estimates of the extremal index based on simulations of length 1,000,000, obtained by the same method as for the rainfall data above. The estimates of the extremal index indicate that the transformed latent Markov chain model does not accurately capture the extremal dependence structure in this example. In particular, the Warren process model appears to underestimate the extremal dependence (i.e. θ is too low), whereas the opposite is the case for the GL-P model. Figure 10 shows estimates of the conditional tail dependence coefficient for the empirical, G-LP and Warren processes, together with the theoretical value for the latent trawl process. We see that the latent trawl and G-LP values mostly agree with the empirical estimate, whereas the WP model underestimates the dependence for this coefficient as well.

	ξ	σ	ρ	κ
Latent trawl	-0.11	20.73	0.17	32.69
G-LP	-0.04	21.09	0.56	32.19
WP	-0.11	20.74	0.96	32.44

Table 2: Estimated parameters for Ozone data using transformed versions of latent trawl model and latent Markov chain model with G-LP and WP chains.

These results show that when the marginal parameters α, β are fixed in the latent layer, the latent Markov process model has less flexibility in the dependence structure than the latent trawl process model. This indicates that in the original hierarchical structure, the marginal parameters also have a strong influence on the dependence structure, and this contributes to the flexibility of the model. The transformed model has the advantage of parameters that are

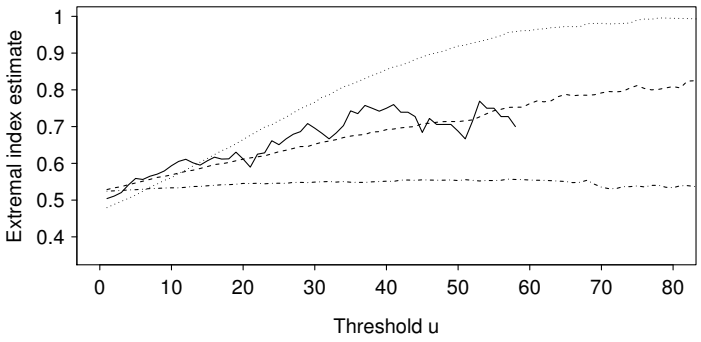


Figure 9: Estimated extremal index for empirical Ozone data (solid), latent trawl (dashed), G-LP (dash-dotted) and WP (dotted).

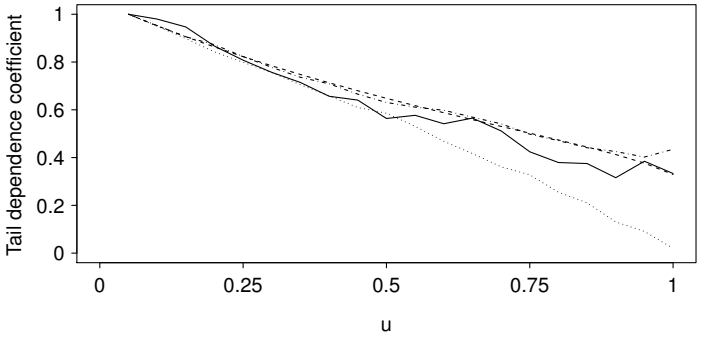


Figure 10: Estimates and theoretical value of the conditional tail dependence coefficient $\varphi(u, u)$ at lag 1, based on empirical Ozone data (solid), latent trawl (dashed), G-LP (dash-dotted) and WP (dotted).

clearly interpretable as contributing to either the dependence or the marginal distribution. In our experience, the parameter estimation procedure for the transformed model also appears to be more reliable, which is not too surprising as the modification of the original model is a form of parameter orthogonalisation.

5 Conclusion

In this paper we have investigated a new model for time series of extremes based on trawl processes. The first part of the paper focused on trawl processes in the context of statistical models, showing how they can be simulated and a suitable process can be chosen to obtain the desired autocovariance structure. Based on these results we constructed an extreme value model that uses the trawl process framework to obtain a flexible dependence structure. This was done by replacing the latent Markov chain in the setup of Bortot and Gaetan (2014) with a trawl process. In contrast with other hierarchical models, this construction has the advantage of preserving the generalised Pareto distribution for the marginals, which is consistent with extreme value theory.

We have also considered a modification to the original model structure that extends the parameter space, allowing for negative shape parameters. To evaluate the extremal dependence we have also developed an adapted version of the tail-dependence coefficient, which can be evaluated analytically for the trawl process model.

The original and modified models were used to analyse two environmental time series, and

compared with the latent Markov chain models of Bortot and Gaetan (2014). For the application using the original model, the results were very similar in terms of capturing the extremal dependence of the data. The advantage of using the latent trawl process was clearer when using the transformed model, where the trawl-based model performed better due to the added flexibility in the latent dependence structure.

There are several aspects of the latent trawl model that could benefit from further investigation. For example, we have used a simple exponential trawl throughout this paper; it would be interesting to look at the result of using different parameterisations for the trawl set. Another possibility is to consider a model where the threshold u is allowed to vary, and then letting the trawl set dependend on u, which should result in a wider range of dependence levels across thresholds.

Acknowledgments

R. C. Noven gratefully acknowledges financial support from the Grantham Institute for Climate Change, Imperial College London. A. E. D. Veraart acknowledges financial support by a Marie Curie FP7 Integration Grant within the 7th European Union Framework Programme. We thank the UK Meteorological Office, the British Atmospheric Data Centre and the London Air Quality Network for providing the data used.

6 Appendix

Latent trawl model pairwise likelihood

The pairwise likelihood can be derived by using lemma 2.16 in the likelihood expression. Consider two observations x_1, x_2 corresponding to time points t_1, t_2. Then we can write down the explicit partition of the trawl sets as follows:

$$\{B_1 = A_{t_1} \setminus A_{t_2},\ B_2 = A_{t_1} \cap A_{t_2},\ B_3 = A_{t_1} \cap A_{t_2}\}.$$

The likelihood of x_1, x_2 is given below for each of the four possible cases:

1. No exceedances, so $X_1 = X_2 = 0$:

$$\begin{aligned} f_{X_1,X_2}(x_1,x_2) &= E_{\Lambda_1,\Lambda_2}\left[(1-e^{-\kappa\rho_1})(1-e^{-\kappa\rho_2})\right] \\ &= 1 - 2\left(1+\frac{\kappa}{\beta}\right)^{-\alpha} + \left(1+\frac{\kappa}{\beta}\right)^{-\alpha\rho|B_1|-\alpha\rho|B_3|}\left(1+\frac{2\kappa}{\beta}\right)^{-\alpha\rho|B_2|} \end{aligned}$$

   where A is the area of the trawl set, and we have used the substitution $u_i \to iu_i$ in the expression for the characteristic function to obtain $E[e^{-u_i\Lambda_i}]$, which is well defined for all arguments $u_i > 0$.

2. One exceedance, we consider $X_1 > 0, X_2 = 0$, the case $X_1 = 0, X_2 > 0$ is similar. We have

$$\begin{aligned} & f_{X_1,X_2}(x_1,x_2) = E_{\Lambda_1,\Lambda_2}\left[e^{-\kappa\rho_1}\Lambda_1 e^{-\rho_1 x_1}\left(1-e^{-\kappa\rho_2}\right)\right] \\ =& E_{\Lambda_1,\Lambda_2}\left[\Lambda_1 e^{-(\kappa+x_1)\Lambda_1}\right] - E_{\Lambda_1,\Lambda_2}\left[\Lambda_1 e^{-\Lambda_1(\kappa+x_1)-\kappa\Lambda_2}\right] \\ =& -\frac{\partial}{\partial u_1}E_{\Lambda_1,\Lambda_2}\left[e^{-(u_1\Lambda_1+u_2\Lambda_2)}\right]\Big|_{u_1=\kappa+x_1,u_2=0} + \frac{\partial}{\partial u_1}E_{\Lambda_1,\Lambda_2}\left[e^{-(u_1\Lambda_1+u_2\Lambda_2)}\right]\Big|_{u_1=\kappa+x_1,u_2=\kappa}. \end{aligned}$$

We have that

$$\begin{aligned}
&\frac{\partial}{\partial u_1} E_{\Lambda_1,\Lambda_2}\left[e^{-(u_1\Lambda_1+u_2\Lambda_2)}\right] \\
&= -\frac{\alpha\rho|B_1|}{\beta}\left(1+\frac{u_1}{\beta}\right)^{-\alpha\rho|B_1|-1}\left(1+\frac{u_1+u_2}{\beta}\right)^{-\alpha\rho|B_2|}\left(1+\frac{u_2}{\beta}\right)^{-\alpha\rho|B_3|} \\
&- \left(1+\frac{u_1}{\beta}\right)^{-\alpha\rho|B_1|}\frac{\alpha\rho|B_2|}{\beta}\left(1+\frac{u_1+u_2}{\beta}\right)^{-\alpha\rho|B_2|-1}\left(1+\frac{u_2}{\beta}\right)^{-\alpha\rho|B_3|}
\end{aligned}$$

Evaluating, simplifying and using that $|B_1| + |B_2| = |A|$, this gives

$$\begin{aligned}
f_{X_1,X_2}(x_1,x_2) &= \frac{\alpha\rho|A|}{\beta}\left(1+\frac{\kappa+x_1}{\beta}\right)^{-\alpha\rho|A|-1} \\
&- \frac{\alpha\rho}{\beta}\left(1+\frac{\kappa+x_1}{\beta}\right)^{-\alpha\rho|B_1|-1}\left(1+\frac{2\kappa+x_1}{\beta}\right)^{-\alpha\rho|B_2|-1}\left(1+\frac{\kappa}{\beta}\right)^{-\alpha\rho|B_3|}\left(|A|\left(1+\frac{\kappa+x_1}{\beta}\right)+\frac{|B_1|\kappa}{\beta}\right).
\end{aligned}$$

3. Two exceedances, so $X_1 = X_2 = 1$:

$$\begin{aligned}
&E_{\Lambda_1,\Lambda_2}\left[\Lambda_1 e^{-(\kappa+x_1)\Lambda_1}\Lambda_2 e^{-(\kappa+x_2)\Lambda_2}\right] = \frac{\partial}{\partial u_1}\frac{\partial}{\partial u_2} E_{\Lambda_1,\Lambda_2}\left[e^{-u_1\Lambda_1-u_2\Lambda_2}\right]\big|_{u_1=\kappa+x_1,u_2=\kappa+x_2} \\
&= \frac{\alpha^2\rho^2}{\beta^2}\left(1+\frac{\kappa+x_1}{\beta}\right)^{-\alpha\rho|B_1|-1}\left(1+\frac{2\kappa+x_1+x_2}{\beta}\right)^{-\alpha\rho|B_2|-1}\left(1+\frac{\kappa+x_2}{\beta}\right)^{-\alpha\rho|B_3|-1} \\
&\times\Bigg[|B_1||B_2|\left(1+\frac{2\kappa+x_1+x_2}{\beta}\right)\left(1+\frac{\kappa+x_2}{\beta}\right)+|B_1||B_3|\left(1+\frac{2\kappa+x_1+x_2}{\beta}\right)^2 \\
&+|B_2|(|B_2|+1/(\alpha\rho))\left(1+\frac{\kappa+x_1}{\beta}\right)\left(1+\frac{\kappa+x_2}{\beta}\right)+|B_2|B_3|\left(1+\frac{\kappa+x_1}{\beta}\right)\left(1+\frac{2\kappa+x_1+x_2}{\beta}\right)\Bigg].
\end{aligned}$$

References

Ancona-Navarrete, M. A. and Tawn, J. A. (2000), 'A comparison of methods for estimating the extremal index', *Extremes* **3**(1), 5–38.

Barndorff-Nielsen, O. E. (2001), 'Superposition of Ornstein–Uhlenbeck type processes', *Theory of Probability & Its Applications* **45**(2), 175–194.

Barndorff-Nielsen, O. E. (2011), 'Stationary infinitely divisible processes', *Brazilian Journal of Probability and Statistics* **25**(3), 294–322.

Barndorff-Nielsen, O. E., Benth, F. E. and Veraart, A. E. D. (2015), 'Recent advances in ambit stochastics with a view towards tempo-spatial stochastic volatility/intermittency', *Banach Center Publications* **104**, 25–60.

Barndorff-Nielsen, O. E., Lunde, A., Shephard, N. and Veraart, A. E. D. (2014), 'Integer-valued Trawl Processes: A Class of Stationary Infinitely Divisible Processes', *Scandinavian Journal of Statistics* **41**(3), 693–724.

Barndorff-Nielsen, O. E. and Schmiegel, J. (2007), Ambit processes: with applications to turbulence and cancer growth, *in* F. E. Benth, G. Di Nunno, T. Lindstrøm, B. Øksendal and T. Zhang, eds, 'Stochastic Analysis and Applications: The Abel Symposium 2005.', Springer, Heidelberg, pp. 93–124.

Bennedsen, M., Lunde, A. and Pakkanen, M. S. (2014), 'Discretization of Levy semistationary processes with application to estimation', *arXiv:1407.2754* .

Bortot, P. and Gaetan, C. (2014), 'A Latent Process Model for Temporal Extremes', *Scandinavian Journal of Statistics* **41**(3), 606–621.

Bortot, P. and Tawn, J. A. (1998), 'Models for the extremes of Markov chains', *Biometrika* **85**(4), 851–867.

Coles, S. G. and Tawn, J. A. (1991), 'Modelling extreme multivariate events', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **53**(2), 377–392.

Coles, S., Heffernan, J. and Tawn, J. A. (1999), 'Dependence Measures for Extreme Value Analyses', *Extremes* **2**(4), 339–365.

Cox, D. R. and Reid, N. (2004), 'A note on pseudolikelihood constructed from marginal densities', *Biometrika* **91**(3), 729–737.

Davison, A. C., Padoan, S. A. and Ribatet, M. (2012), 'Statistical Modeling of Spatial Extremes', *Statistical Science* **27**(2), 161–186.

Davison, A. C. and Smith, R. (1990), 'Models for exceedances over high thresholds', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **52**(3), 393–442.

Embrechts, P., Koch, E. and Robert, C. (2015), 'Space-time max-stable models with spectral separability', *arXiv:1507.07750* .

Embrechts, P., Mikosch, T. and Klüppelberg, C. (1997), *Modelling extremal events for insurance and finance*, Springer Verlag, Berlin.

Environmental Research Group, King's College London (2015), 'London Air Quality Network database'. Available from `http://www.londonair.org.uk`.

Ferro, C. A. T. and Segers, J. (2003), 'Inference for clusters of extreme values', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**(2), 545–556.

Fuchs, F. and Stelzer, R. (2013), 'Mixing conditions for multivariate infinitely divisible processes with an application to mixed moving averages and the supOU stochastic volatility model', *ESAIM: Probability and Statistics* **17**, 455–471.

de Haan, L. and de Ronde, J. (1998), 'Sea and Wind: Multivariate Extremes at Work', *Extremes* **1**(1), 7–45.

Hedevang, E. and Schmiegel, J. (2014), 'A Lévy Based Approach to Random Vector Fields: With a View Towards Turbulence', *International Journal of Nonlinear Sciences and Numerical Simulation* **15**(7-8), 411–435.

Heffernan, J. E. and Tawn, J. A. (2004), 'A conditional approach for multivariate extreme values', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **66**(3), 497–546.

Hsing, T., Hüsler, J. and Leadbetter, M. (1988), 'On the exceedance point process for a stationary sequence', *Probability Theory and Related Fields* **78**(1), 97–112.

Huser, R. and Davison, A. C. (2014), 'Space-time modelling of extreme events', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(2), 439–461.

Kallenberg, O. (1976), *Random Measures*, Academic Press Inc. (London) Ltd.

Leadbetter, M. R., Lindgren, G. and Rootzén, H. (1983), *Extremes and Related Properties of Random Sequences and Processes*, Springer Verlag, New York.

Ledford, A. W. and Tawn, J. A. (1996), 'Statistics for near independence in multivariate extreme values', *Biometrika* **83**(1), 169–187.

Ledford, A. W. and Tawn, J. A. (1997), 'Modelling dependence within joint tail regions.', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **59**(2), 475–499.

Ledford, A. W. and Tawn, J. A. (2003), 'Diagnostics for dependence within time series extremes', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**(2), 521–543.

O'Brien, G. L. (1987), 'Extreme Values for Stationary and Markov Sequences', *The Annals of Probability* **15**(1), 281–291.

Pickands, J. (1975), 'Statistical inference using extreme order statistics', *the Annals of Statistics* **3**(1), 119–131.

Rajput, B. S. and Rosinski, J. (1989), 'Spectral representations of infinitely divisible processes', *Probability Theory and Related Fields* **82**(3), 451–487.

Sato, K. I. (1999), *Lévy Processes and Infinitely Divisible Distributions*, Cambridge University Press.

Smith, R. L., Tawn, J. A. and Coles, S. G. (1997), 'Markov chain models for threshold exceedances', *Biometrika* **84**(2), 249–268.

UK Meteorological Office (2012), 'Met Office Integrated Data Archive System (MIDAS) Land and Marine Surface Stations Data (1853-current)'. NCAS British Atmospheric Data Centre. Available from `http://badc.nerc.ac.uk/view/badc.nerc.ac.uk__ATOM__dataent_ukmo-midas`.

Varin, C. (2008), 'On composite marginal likelihoods', *AStA Advances in Statistical Analysis* **92**(1), 1–28.

Varin, C., Reid, N. and Firth, D. (2011), 'An overview of composite likelihood methods', *Statistica Sinica* **21**, 5–42.